

URL: <https://ecorner.stanford.edu/videos/responsible-ai-innovation-entire-talk/>

Rahul Roy-Chowdhury is Grammarly's CEO as of May 1, 2023. He was previously the company's Global Head of Product, overseeing everything that makes the company's writing assistance software tick, managing the product, design, and data science teams. Before coming aboard at Grammarly, Roy-Chowdhury spent many years at Google leading the Privacy, Safety, and Security teams to unified solutions across Google's product portfolio. In this presentation and conversation with Stanford adjunct lecturer Ravi Belani, Roy-Chowdhury Shares his company's approach to building ethical AI tools and encourages aspiring entrepreneurs and tech leaders to get engaged in the conversation about responsible AI.



Transcript

Host Who you are defines how you build. 00:00:06,180 (upbeat music) - Welcome Stanford and YouTube community 00:00:11,010 to the Entrepreneurial Thought Leader Seminar at Stanford University, Brought to you by BASES, the Business Association of Stanford Entrepreneurial Students and STVP, the Entrepreneurship Center in the School of Engineering at Stanford. I am Ravi Belani, a lecturer in the management science and engineering department at Stanford, and the director of Alchemist, an accelerator for enterprise startups. Today we are thrilled to have Rahul Roy-Chowdhury, CEO of Grammarly, here with us. How many people have used Grammarly? Oh wow, okay. So it almost needs no introduction. How many people have used Chrome? - Chrome? 00:00:50,310 and many more, the man behind them is our guest speaker tonight, Rahul. So Rahul is the current CEO of Grammarly, and used to be the global head of product of Grammarly. Is coming to the CEO position from a deep product background. And Grammarly, as you guys all know, and it needs no introduction, is an AI and NLP driven writing enhancement tool whose mission is to make all lives better by improving communication across people.

Rahul formerly spent 14 years at Google where he really honed his chops in product management and he rose up the ranks to become ultimately the vice president of product at Google. And on the way he was overseeing products like Chrome and Chromebooks and many, many others. He's a grad of Stanford's from the business school. He got his MBA from the Stanford Business School. and before that he was also getting a PhD in AI. So he was, he's been immersed in computer science and AI before AI was cool, when some times people called it AAI, artificial artificial intelligence way back when. And now it's become this reality that we see. But Rahul was originally getting a PhD from Columbia, and he dropped out with a masters because he was called to have impact in the world and to leave, you know, the safe realms of academia to actually create product. Before that, he also got a bachelor's degree in math from Hamilton College. And so today we're gonna have a discussion about all things, but centered also around AI and this powerful wave of AI that's ushering into all of you.

And then we're gonna have an interactive conversation. So prepare or think about all of your questions, everything that's fair game. And please join me in welcoming Rahul to ETL. (audience applauds) - Thank you. 00:02:28,680 Thanks Ravi, appreciate it. Hey everyone. Great to be back on campus. Wow, this is so awesome. I wanna talk to you about the future of AI. But before I do that, let me ask you all a really, really important question.

What do you call an AI that has learned how to take a picture of itself? Anyone? You call it selfie aware? Get it? (laughs) I'm famous, or I should say I'm infamous in Grammarly for sharing terrible jokes. I just had to get one in. And, you know, the future of AI, that's a big, that's a big topic. And so it's almost preposterous. Who am I, why am I talking about this? What does this mean? And what I'd like to share with you all today is we at Grammarly have been building and using AI to help our users for a long time. Many of you are users, thank you. And we have developed tools and technologies and an understanding of what it takes to build an AI system at scale. And I wanna share what we've learned in the hopes that you can get benefit from it. Some of this can be broadly applied, it's not specific to Grammarly. But really what I want is I want this to be a conversation.

I want the end of my session to be the beginning of a conversation, an important conversation. A conversation about how AI can be used to help us, not harm us. How AI can be used to increase our capabilities and our potential, not reduce it. How AI can be used to augment our skills, to help us communicate in a new and better way, not rob us of our voices. So that's the conversation I'd love to have with you all today. So to tee up the conversation, let's kind of just do a situational assessment. Where are we? Now the reality is AI is the most powerful technological force that is shaping our world today. No question in my mind at least. It is already, not future, already having a massive impact on almost any industry that operates at scale. Think of manufacturing.

AI is being used to create better components. AI is being used to optimize global supply chains. Think of healthcare. AI is being used to help radiologists detect cancers earlier, faster, more accurately. Think of medicine. AI is being used to help in the drug discovery process, to analyze billions of compounds to pinpoint the most promising areas to focus research on. Think of software. AI is enabling new and personalized tools to help us be more productive in our professional and personal lives. And we are just getting started. In our lifetimes, AI is gonna help us solve the most important problems, things like global warming.

And with the rise of generative AI, which I'm sure is on everyone's mind, AI is gonna help us with creative expression as well, to take a high level idea and turn into something concrete: a work of art, a piece of writing. And that is awesome. I am at heart, I've been working in technology for a long time, and at heart I'm a technology optimist. I believe that technology is gonna help us make our lives better, improve the world. And I've seen this happen over and over again in my career. I remember, we talked about Chrome, in the early days I remember when I worked on Chrome, I saw the impact that access to the web and low price smartphones, bringing access to information to billions of people around the world, and I saw the impact that had. That was technology, helping people at scale. So AI is fantastic, it's awesome. But with great power comes great responsibility. And AI can also cause real harms either by design or just as an accidental byproduct.

Not intentional, but it just happened accidentally as a consequence of the technology being very powerful. So you can think of things like AI-based recommendation engines. They've caused real harm and they've led to societal problems. Things like radicalization or addiction or just getting so immersed in this technology that you lose interest or don't spend time on real world interpersonal relationships. I mean, our brains are no match for these precisely engineered hits of dopamine that these systems can deliver to us. And as a parent of young kids, I worry about this. I mean, I want my kids to get access to all the great things that AI has to offer, but I wanna protect them from these harms. And if you think about AI making decisions for consequential things in our lives, things like screening job applications, or making decisions about who gets approved for a loan and who doesn't, those systems have bias in the underlying data and those biases cause real harm. And the fact is, people are getting better, systems are getting better, things are improving, so it's not like this is a gloom and doom story, but the harm is real. If you, I don't know if you, if you watched John Oliver's "Last Week Tonight" talk show.

It's really funny, it's awesome, it's a great show. And a couple of months ago his show was, he did a show on AI and it's well worth watching, especially given this class and this talk. But one of the things he highlights is AI being used to screen job applications. And you know, they go through a lot of issues and bias and how the system has a lot of challenges. But the one thing that really struck me from that John Oliver segment was one of these leaders of these companies providing these services, his advice to job seekers is, hey, don't stand out. Your job is not to be unique or special or differentiate yourself. Your job is to get past the AI gatekeeper so you can actually get to a real human. And that's awful. That's so disempowering. That's so dystopian.

That's not a future I want for myself. I'm sure that's not a future you want for yourselves. And so the harm is real. So the question is, what do we do? There's real benefit. This is a very powerful technology, but there's real harm. And so what I see is people pick a couple of lanes, one lane people pick is they say, oh my goodness, stop. Let's stop. It's too dangerous, too difficult, we can't move forward. That's not great because the benefits are very real and we should get those benefits. I've seen another lane people pick where they say, let's just go full steam ahead, innovation at all costs.

We'll hope for the best. And that's not great either because the harms are real and hope is not a strategy. So we don't wanna stop, we don't wanna just heedlessly charge ahead. What should we do? Well, I think we should pick an intentional path forward. And the thing I really want to get across to all of you is we decide, all of us in this room, folks on livestream, we get to decide how this technology gets developed and deployed. Don't let the technology be done to us, let us decide how this technology is deployed and how it affects us in our daily lives. We have the power, don't cede that power. And that's a really important thing for you all to remember, now, as you go into the job market, as you embark on your careers, just remember

that we have the power and the control to decide how this technology is shaped and what shape it takes in our lives. And that's the conversation that really today is about. So at Grammarly, we are working on frameworks and technologies and tools to try and bring AI to people at scale.

And we have the basic belief that AI is about augmenting human potential. We call it augmented intelligence. And you can think of augmented intelligence as the belief that AI is successful when it augments us, when it makes us better, when it enhances our capabilities. And that should be the goal of these systems. In fact, I really wish we could have a do-over and rename AI to stand for augmented intelligence instead of artificial intelligence. Because I actually think that that is really what we want out of these systems. We want them to help us. Artificial intelligence leads you down these conversations about is it conscious, is it sentient? Really, we want these systems to be utilitarian and help us. And so we wanna build systems and we wanna deploy systems to our users that remain true to this vision of augmented intelligence. To do that, we use something we call the TRUE framework.

The T is trust, which is a commitment to world best in class privacy and security. The R is responsibility, which is a focus on fairness and user safety. The U is user control, which is really putting users in control and making sure users have agency through the entire experience. And the E is empathy, which is making sure that we are actually solving, we understand first of all what our users need and we're solving a real user problem. So that's a true framework. And what I'd like to do is give you a case study of how the TRUE framework gets applied. Just real practical advice based on building systems at scale. So do do that, let me take a step back and let me explain to you what Grammarly does because that'll help you understand the case study. So Grammarly's in the business of helping people communicate better. There we go.

And the thing you should know about communication is it is just really hard to get right. We all have different styles, we have different contexts. We communicate using lots of different tools. And the chances of someone being understood exactly the way they intended using the right tone, using the right information is not that high. And in fact, the number of tools we're using to communicate has gone up because the tools landscape is fragmented. And the way we work has changed dramatically in the last couple of years. There's much more remote work, distributed teams, global teams trying to work together with different contexts. So communication has always been hard to get right. It's even harder now than it was before. And the cost of not communicating very well is very high.

In fact, we did a study that said that the cost of ineffective communication just in the U.S. in a single year is \$1.2 trillion. And this number, this was about a year ago, this number is not going down. It's going up because the volume of communication is going up and the effectiveness of communication is going down. So this is a hard problem to get right, but when we do get it right, it's fantastic. We are able to get our point across we feel like we understand each other, we feel like in a company we can collaborate better, get things done. And companies report this, they see better outcomes. They see teams are in flow state, they move faster, they achieve better outcomes. So that's the communication problem. That's the problem we are solving at Grammarly.

Now everyone knows, yeah, generative AI is the thing, everyone's working on generative AI. We've been working on generative AI at Grammarly, and about a month ago we launched our take on generative AI, a product we call Grammarly Go. And so to bring Grammarly Go to our users, we have millions of users around the world, we use the TRUE framework. So let me walk you through it in detail, how do we think about this framework when we thought about this product that we just recently launched to our users? So let's start with trust. Trust is about making sure that you understand what user's expectations are around their data. These are questions like, well, what data are you collecting? These are questions like, well, what are you doing with the data? How long are you keeping it for? And really at the root of all of these questions, I think is an underlying question of, hey, what are your incentives? And are your incentives and my incentives, are they aligned or not? And so at Grammarly, we have built a business and a company around this idea of making sure that our incentives and our users' incentives are always in 100% alignment. And so it's a very simple decision making process. A new business model comes our way, and we think about it and we say, okay, there's revenue potential that's interesting, but does it cause us to lose alignment with our user's interests in the way we are thinking about user data? And if the answer to that is, yes, we're not gonna do it. So user alignment is critical. That's really the root of privacy is understanding alignment.

If privacy is about alignment and what are you gonna do with user data, security is about keeping that data safe from harm, from misuse. And at Grammarly, we've invested a lot in security over the years. We have the whole alphabet soup of certifications, SOC2, HIPAA, CCPA, you name it, we've got it. And these things don't happen quickly. They take a long time and take a lot of investment to do, and it's hard to bolt on after the fact. And now with new technologies like large language models, new threat vectors emerge and the full landscape of threat models for something like a large language model isn't fully understood. And so the best way to approach that uncertainty is to really have a multi-layered security story. You're not just relying on one thing to work well, you're relying on multiple things working together, having this portfolio approach, things like a bug bounty program, or things like having a red team. The concept of a red team is you have a team, we have this team at Grammarly, and their job is to proactively find vulnerabilities in your systems so that we find them and fix them before an outside hacker does. And so really investing in this portfolio of security approaches is the right way to deal with a new piece of technology.

The thing I just wanna kind of give you as a takeaway for this whole trust idea is, it is very hard to earn, it's very easy to

lose, and it's very hard to bolt on after the fact. So as you embark on your careers and you think about all the cool apps and businesses and technologies you're gonna build, keep trust in mind because trust needs to be the foundation of everything else you do because we need to earn and strengthen that trust as the foundation. It's gonna be very hard to scale a successful company without that basic foundation in place. So that's trust. So with Grammarly Go, we applied this framework, these technologies, all of our investments, they helped us launch Grammarly Go, keeping our trust posture intact with our users. Responsibility is about making sure that we are focused on fairness, we are focused on user safety. And there's a lot of technical work to be done here. Things like evaluating and making sure you have very high quality training data for your ML models, making sure that data is free from bias. Making sure you can do things like track and find sensitive text, like hate speech. And so for example, one of the things we've invested in at Grammarly is a sensitive text classifier.

And so you can think of a large language model generating text, and when our classifier catches that text and says, "hey, that's sensitive text," we can then take some appropriate action. For example, we can say, well, if a user sees this text, it's not gonna be a great experience for them, so let's just not show them this text. So we can actually keep our users safe from the underlying harm that these models can cause. And so we invest in these technologies and I would highly encourage everyone thinking about building AI at scale to really think deeply about how to do this well, quality, evaluation, sensitive text, classifiers to detect these things. These are important things to get right, but it's not just about the tech, because no matter what you do, no matter what you build, no matter how good your classifiers are, there's always gonna be some unexpected thing that you're gonna run into when your users encounter your product. And what are you gonna do then? And that's a question you gotta think through ahead of time and not react after the fact. So at Grammarly, for example, the way we approach this is we give users a lot of different ways to send us feedback: in the product, email our customer service team. And it's not just giving us feedback. We have a process so that when we get that feedback, we can take action. Someone is looking at that and saying, "oh yeah, this is a problem.

We gotta fix this." Who's gonna fix it? Who's on call? Let's fix it. Let's roll out the fix. Let's keep our users safe. So a key part of responsibility is not just the technical innovation that's happening, which is important and great, but also making sure that you have some plan for the inevitable unexpected thing that's gonna happen when your real users, use your product in their real lives. And so what's your answer then? And it's important to have an answer. It may not be exactly the Grammarly answer, but we need some answer. So that's responsibility. User control is really the idea that users should feel like they have control over the entire experience. So they don't feel like like the technology is being done to them. They feel like they are in control of the experience, they are driving the process and they are driving the outcomes.

So with Grammarly Go for example, we are generating text on behalf of our users. So number one, we wanna make sure users feel like the text generated by Grammarly Go sounds like them, represents their voice. So we built a feature called My Voice where users can control how they sound. So it's not just boiler plate, generic text, it reflects how you wanna sound. And then the entire experience allows users to see what's happening. We provide nudges, we provide prompts, but ultimately we always defer to the user to make the decisions. And that's not an accident. That's not just a thing that happened. That is a deliberate design choice. And so as you're designing and building these features and products, you're gonna be faced with these choices.

Do I bring the user in, do I not? And my suggestion and my hard one experience tells me, bring the user in, make sure that they feel they have control over the experience. That's the user control piece of TRUE. And then the final piece is empathy. And empathy is the idea that we are trying to solve a real user problem. And it sounds like an obvious thing, like of course that's what you would do. Why wouldn't you do that? But in fact, what I have seen over and over again in my career is that it's very easy to lose sight of that because there's so many distractions. You're busy working on your product and feature, you started out with the best of intentions, here's the user problem, I'm gonna solve it. And then stuff happens like, oh man, my competitor over here did this other thing, and so I gotta respond. Or here's this cool new tech, I gotta get it out there. Let's see what happens.

And so it's very easy to get sidetracked and you're kind of looking at a bunch of different things and all the time you spend doing those things, it's time you spent not thinking about the user. And when that happens, usually your product is going astray in actually helping solve a real user need. And so you've really gotta will yourself, sometimes unnaturally, to come back and ask yourself the question of what is the user need? What is the problem we are solving for them? And how am I confident that my solution is actually a solution to the problem they actually have, not just something I've come up with because it's pretty cool and I'm just gonna roll it out there and see what happens. So empathy is something that is really important to keep bringing back to the conversation as you build and deploy these systems. So that's a true framework. Now we use it at Grammarly, but there's nothing specific to Grammarly about this. So my hope is that you all can see some value in this and apply this in your jobs, in your future jobs, in your careers, as you think about building applications using this framework. And imagine a world in which magically the TRUE framework became the defacto way of doing things. Imagine a world in which all apps, all industries, everyone was using this TRUE framework. What would that world look like? Let's just kind of try to go through this visualization exercise.

Well, that's a world in which people are gonna be able to feel that they are empowered to do better, to achieve better outcomes, that they're better understood. That's a world in which people will feel a deeper sense of connection with each other, a deeper sense of belonging. They'll work better, businesses will be more productive and will achieve better outcomes.

And that's a world in which we are deploying AI at scale to help people achieve better outcomes, help businesses achieve better outcomes. And that sounds like an awesome world. That's the world I want. That's how we get the promise of AI that we talked about in the beginning without all of the issues that come along with it, if we don't think about it methodically and intentionally. So remember, we have the power, we have the control. We decide how this technology gets deployed and built. You all are leaders in tech or future leaders in tech.

So my wish and my hope is that all of you think deeply about how you can use that power. And my request is please commit to using that power to build AI that augments human potential, that makes humans do better work, to improve our potential, our capabilities, to give us superpowers. I mean, if we don't make that commitment, who is going to? So let us make that commitment because we all want to build a better future using AI. So let's make that true. Thank you. (audience applauds) - Let me ask one or two questions 00:26:04,500 before we go to the students. First of all, thank you, Rahul, for providing that framework. It's fantastic to actually have something actionable when thinking about generative AI. - Yeah, excellent. 00:26:15,080 - Not everybody is as empathetic as, you know, 00:26:16,590 as Grammarly is.

Do you think that having frameworks themselves are sufficient? Or what's your position on government regulation? Do you think government regulation, do you invite that or disinvite that? - Yeah, I think 00:26:29,400 my take on government regulation is it's an important conversation. As I said in the beginning, it's an important conversation. So I'm very glad that the conversation is happening, first and foremost, because it's critically important. I think government regulation for technology this powerful is needed and welcome. But really, I would say that all of us as leaders in the technology field, as future leaders in the technology field, we should be taking an approach of responsibility that is forward looking. So if I'm doing everything right, when there is regulation and there's clear guidance, I'm already doing those things and more because I've actually thought deeply about how I deploy these systems at scale responsibly, not because of regulation, because it's the right thing to do for our users. - And do you think the right thing 00:27:16,620 will also win in the marketplace, or do you think if there is a ruthless competitor that has no morals, will they win the generative AI space over a ethical counterpart? - I believe so because I've always seen 00:27:35,377 that technology creates dilemmas, but technology always finds a way out of those dilemmas. And technology always creates good outcomes, ultimately becomes a force for good. And I've gotta believe that that's gonna be true for AI. But it's not a law.

But the thing I wanna make sure we all just understand is this is not a priori thing. This is not a law of nature. This is us, our behavior and our actions bring about these outcomes. So the question of like, you know, do we make sure the good guys win or not? It's not that there's a law of physics that'll enable that outcome. It's us, it's our collective actions that'll make that true or not. And so this is why I'm actually excited to have this conversation with all of you, because I just wanna make sure you understand the power you do have to shape that outcome. - And I'm gonna ask one more question 00:28:23,227 and then we'll open it up. But is it pernicious, because we had Open AI's Chief Science Officer, Ilya here a couple weeks ago. - Yeah. 00:28:33,480 what strategic decisions they're gonna make in the future, this is mainly Sam, his co-founder, he'd oftentimes answer by saying, well, we're gonna ask the AI, you know, what to do.

And the AI is so integrated now into the processes themselves that this line of where I, as a human end and the AI begins can get very blurred. Just to your example of the feed, you know, influencing even your senses of alignment and incentives can be affected by the generative AI itself. Can we stay distinct from the generative AI if it's so deeply integrated into the product development process? - I think we can 00:29:05,370 and I think we should. - Okay. 00:29:06,240 - I think we absolutely should. 00:29:08,190 This is about user control. This is a good example of let's use the AI to help us. Maybe AI can help us create great ideas and figure out, well, what do I do next? What's my product roadmap look like? It'll be great if the AI helps me figure out Grammarly's company strategy, but ultimately it has to be us. We have to be in control and we should not give up that control. And that's important to build into our systems.

- Okay. 00:29:33,510 Audience Member Hey Rahul, awesome talk. 00:29:36,090 I'm actually a PM at Google working on generative AI. - No kidding. 00:29:38,070 Bards, love it. Audience Member One thing that kind of bothering me, 00:29:41,220 not bothering me, kind of keep me up the night is, as a user, I see the potential of AI. I love to using AI to create different things. Like let's say I write a song just like Beyonce. - Yep. 00:29:57,420 the news industry, the music industry.

So how do we balance kind of the AI capability from the assistant right now to become the actual content creator? And how do we give the power back to different industry? - Yeah, yeah, I think there is a lot of uncertainty 00:30:13,260 and a lot of panic. I think that's true of every big tech innovation. If you go back in time in the '70s was the first time that ATMs were deployed at scale in the U.S. and bank tellers panicked because they're like, oh my god, ATMs, our jobs are done. And actually there was vandalism happened. ATM machines were vandalized because people were just fearful. What ended up happening is that because of the widespread use of ATMs, it was much cheaper to open a bank branch because you didn't need as many people to staff it. And so bank tellers could move on to do more meaningful, more complicated, higher order tasks. And the number of bank tellers employed in the U.S. pre-ATM was lower than the number of bank tellers employed in the U.S.

post-ATM. Technology always finds a way to help us, but we need to make sure that we make that true. So like, you're a PM at Google, the stuff you build really makes a difference to billions of people around the world. So bring these principles to mind. Think about, you know, you're building this thing that is very consequential. Don't throw it over the fence. I'm not saying you're doing it, but I'm just saying be intentional. Put the user first. Make sure that you're solving user problem. Have

empathy for the user task.

Make sure user control is front and center. So we shape this future based on the actions that we collectively take. Each individual action as we added up over hundreds and thousands of engineers and PMs working on really cool stuff. That is what makes the future. - Awesome, thank you. 00:31:50,490 - Yeah. 00:31:55,440 Did you select high quality data or did you work with web crawl services? Or how did you select the data? - Yeah, we, with the launch of Grammarly Go, 00:32:04,620 this is a very fast evolving space. With the launch of Grammarly Go, we used GPT3.5 turbo running on Azure. And so what we did was we applied all of the other deployments we used, such as the sensitive text classifier as a post-processing step. So you can look at the output generated and say, well, this is not great, change it.

Things like My Voice, think of that as prompt engineering into that underlying model. And then in the future, the space of these foundation models is changing very rapidly. So in the future, there may be some other model that's better, maybe it's an open source model that's better. Maybe it's a local model running in a device that's better. And so we've invested a lot in just quality evaluation. So if a new model comes along, we have a team of experts, human experts in place, that can evaluate models for quality, for bias, for fairness, for safety. And so we are ready to make that change as these models evolve. - [Audience Member 2] I have a question about, 00:33:02,400 it's actually a follow up to what Josh was asking about LLMs, right? You had touched, Rahul, on trust and safety, et cetera. There's, as you know, there's been a lot of talk about the inputs that go into LLMs and the vastness of the data and the transparency of the data. As an individual user, could you comment on how me or anybody else can get a sense of what goes in and how as a user I get to have either a say or some sort of visibility into how these LLMs are being trained, and how I can get a sense of sort of control or visibility? - I think that's really, 00:33:42,420 what you're asking is really an artifact of the design of the system as opposed to how the model was trained.

As a user, it's very hard to understand what the training data was that went into generating a specific output. As you said, the space is very large. So as an example of what you can do is, with Grammarly Go, you can create your own, like you can create your voice. You can say, this is how I wanna sound, this is my voice. And so the product experience enables you to control what goes in, what comes out. So my feeling is the way that you solve this user conundrum of there's so much complexity, how do I explain to a user what this complexity is, is really through very thoughtful design, where the user, it goes back to user control. And it's really important that users feel like they understand and have a voice in shaping that experience. - [Audience Member 3] My question is like, 00:34:39,120 there is some extent that AI can do certain things, but there's some, like, it's, right now at least we are still figuring out how can it make the AI to be as close to a human coach as possible to give people like real time feedback and coaching in the context. I'm curious, like from your own experience in research, like what can the AI do in terms of improving communication? What are the things they cannot do yet? And where do you see like how it's gonna evolve over the years? - Yeah, yeah, good question. 00:35:09,090 I, you know, I've used, I use AI, I use Grammarly Go in my day-to-day life.

One of the things that I find very useful is there's a feature called quick reply. So I get an email, I get a ton of email every day, and maybe there's an email from someone saying, "hey, I'd like to meet to discuss some topic." And so what Grammarly Go does is understands the email, understands the intent of the email, and gives me some choices on how to respond. Like, yes, accept, suggest sometimes or politely decline. And it generates an email for me that's kind of a pro forma email, but it does it with full context in my voice and it generates an email that I actually feel happy and proud to send as opposed to like just a generic piece of boilerplate text that sounds robotic. And I use that a lot. And I just, it's a huge productivity hack for me, and it just helps me in my day-to-day work. There are many things that LLMs can't do yet, but honestly the space is changing so quickly, building a model of the world, reasoning skills, they all seem to be emergent properties as these models get bigger. And it seems somewhat, I'm not a researcher in the field, but it does seem a bit of a mystery at what properties emerge from each next generation of model. And I don't know anyone who's really understood that so we'll learn. And you know, the last six months have been insane.

I mean, just think about six months ago, maybe you folks are an exception, but most people had no idea what an LLM was. It's like just suddenly it just took over the world in six months. It's crazy, right? It's happened so quickly. So how this proceeds over the next six months, 12 months, 24 months? I'm not quite sure. I think we are just gonna have to see. - And if you're a Stanford student just passionate about 00:36:56,010 generative AI with all the time in the world, where would you focus your time if you wanted to do something entrepreneurial? - Ooh, hard to, that's interesting. 00:37:05,550 You know, I think, here's what I would say. I would say that focus on a user problem. It's not a thing to say I'm gonna do generative AI. That's my view on it.

Other people may disagree. It is a thing to say, I'm gonna solve a user problem and I can do it really well using generative AI. So, you know, one of the things that I've learned in my career is it's really good if you fall in love with problems, but don't fall in love with solutions. Because the solutions will change over time. And often what ends up happening is we end up falling in love with the solution (laughs) and not the problem. You don't even realize it's happening. And then when there's a better solution out there, you're like, no, but my thing that's just so awesome. It's actually, you know, it's better not to get too attached to any one solution. So I would say really the best advice I can offer you is find those problems, find problems worth solving, start there, and generative AI will be a great way to solve those problems. - Awesome.

00:38:10,263 Can I ask, do you view Grammarly as a product or a platform? - It is, it's a product. 00:38:15,750 - It's a

product. 00:38:22,620 You can think of as it's a product that layers onto every other product that you use. - But even as you're releasing your own model, 00:38:31,920 you're not conceiving of that as this ecosystem that's gonna be built on top of Grammarly? - Yeah, I mean, I think in the fullness of time, 00:38:39,240 we'll move from, you know, today we have multiple products. We have an extension, we have mobile products, a native application. And at some point we'll move from sort of a suite of products to a platform. But right now we are very focused on, we have these entry points to our service, and that service works everywhere that people write. That's what we wanna do. - Terrific. 00:38:58,250 I think we have a question here at the front row.

Yes. - [Audience Member 3] So my, my question is about like 00:39:03,270 every time that we have a technological shift, we get like new ways of communicating. What are you, what do you think are gonna be the ways that are gonna be enabled through this revolution of generative AI? And like, what do you think about this meme of like my AI, when I be talking to your AI or something like that? - I hate that meme. 00:39:22,110 Let's make sure that meme doesn't come true. Because look, I mean, okay, so let me give you some, let me give you some real data points. I talk about the \$1.2 trillion, you know, loss thing, Here's the reality. The average knowledge worker in the U.S., this is a global problem, but just U.S. numbers, average knowledge worker in the U.S. spends more than half their work week, I mean, I should say we, I mean I'm one of them and it's true for me, I spend more than half my work week engaged in written communication. That number has gone up by 18% in the last 12 months.

So that's going up, which means we're writing more. The effectiveness of that written communication during that same 12 month period has gone down by 12%. So the volume is going up, the quality is going down. So we can't outrun this problem. We can't solve this problem by saying, let's just do more. If I can only just run a little bit faster on the treadmill, you know, salvation is at hand on the other side of that. Actually, the more you do, the worse it gets. So really the answer has to be, we've gotta communicate better, not more, but better, because that's what the data says. It's not, I'm not making this up. This is what the data says.

You gotta figure out how to communicate better. So if you think of, if you think about generative AI, we're already communicating so much, right? We're already lowering the quality of communication. And then you think about the cost of communication or cost of content going down. It's just now easy to create reams of blog posts and what have you with generative AI. You can do that. The technology enables you to make that true. But what's that gonna help? You're just gonna increase that 18% to like 50%, 60%, you're gonna drive that 12% down to like minus 50%. That's not gonna help anything. So that's why you really go back to user problems. And if you look at the user problem, you're gonna realize very quickly that it's not more, but better.

And then you, when you think about what business do I start with generative AI or how do I use it? That's your mantra. And that'll guide you to find the right solution. - How do you prioritize, 00:41:27,300 given that everything's changing so rapidly? - Yeah. 00:41:53,160 So there's some things you could figure out, like what are the user problems you're solving? Why are these the most important user problems? What data do we have to back it up? We talk to users, we conduct interviews. So there's a lot of kind of input coming in to help us understand the problem space. Then there's things like, well, what's the privacy story? Are we aligned, are we not aligned? So you can kind of very clearly start using filters there. And then you can do some sort of an impact assessment. Well, okay, this thing would be great to do, but the impact might be marginal. This other thing would take the same amount of time to build, but actually would have massive impact to our users, so let's focus on maximizing impact. So there's a set of these things that you kind of put in place as guide rails and guardrails and filters.

But at the end of the day, there's judgment and there's, you know, there's teams spend many hours inside Grammarly, inside Google, inside every company really debating and discussing and arguing about the roadmaps and priorities. And you know, maybe at some point AI will just tell us what to do. But right now it's really something we have to really engage in a lot of debate and discussion on. - [Audience Member 4] I'm curious about 00:43:04,500 when you were speaking about responsibility and this sensitive text feature. I guess kind of to play devil's advocate, where for you is the line between filtering out, like potentially triggering speech and like censorship and these AI sort of like taking, you know, a limiting, yeah, access to certain ideas? - Yeah, it's a great question 00:43:31,440 and it's a question we ask ourselves every single day. I don't know that there's any perfect line etched in stone or in paint or whatever that metaphor is. We have a team of linguists. They are tracking and studying authoritative sources of what constitutes different kinds of bias, hate speech. And we try to assemble from a set of these sources sort of rules of the road. You know, what is the thinking around this? What's the research, what's the scholarship? And how can you apply that in a meaningful way in the product? And that changes, right? So there's not like you, you do it and you say, okay, check the box, I'm done.

There's a constant evaluation. When COVID became prevalent, one of the things that was new was lots of biased ways to refer to COVID. And that became a thing. And so Grammarly's sensitive text classifier caught those things and was able to suggest to users, I mean, at the end of the day, we can't write for you, but we'll suggest to you that, hey, the thing you're writing is actually considered pretty biased as a way to refer to this virus. And so maybe you should consider this, and here's why, and you should consider this other thing. So this is a thing that keeps changing and adapting and we are trying to be as close to current as we can. Yeah, thank you. - I have to draw 00:44:53,400 this session of ETL to a close. Thank you, Rahul. Thank you so much for that.

- Thank you. 00:44:55,966 This is awesome. Thanks for your time, I really appreciate it. (audience applauds) - Next week we're gonna be closing out the year 00:45:06,060 with Josh Wolfe, who's the managing director of Luxe Capital. So we'll have a venture capitalist in for the final talk. - Wow. 00:45:15,000 on the eCorner YouTube channel or this and other talks are also at ecorner.stanford.edu. So thank you guys for coming to this week's ETL. - Awesome. 00:45:25,049 (upbeat music)..